Notes on the reliability of the HST gyros

Henning Leidecker
Code 562, GSFC
301-286-9180, hleideck@pop500.gsfc.nasa.gov
and
Walter Thomas
Code 302.0, GSFC
301-286-8896, wthomas@pop300.gsfc.nasa.gov

September 10, 2001

Abstract

The most important failure mode of the HST gyros is *flex lead rupture*, caused by corrosion. This mechanism is characterized by a failure rate that is initially vanishingly low, and rises to large values only after a threshold time (determined mainly by the corrosion rate and not the initial strength of the flex leads) is passed. The threshold time is a random variable, and so the failure time is also a random variable. The observed five times-to-fracture and the eleven observed operating times with no fracture, among the sixteen gyros to operate so far, are well fit using a Weibull Probability Law with time-scale $\eta = 5.89$ yr and shape-parameter $\beta = 4.82$: the correlation coefficient is $r = 0.989$. This fit excludes the previously used Exponential Probability Law with better than 95% confidence.

The total Probability Law for the lifetime of an HST gyro is the product of the Weibull Law for corrosion-induced rupture, with an Exponential Probability Law describing the other failure modes (electronic failure and lube patch failure) and an estimate of the latter's mean-time-to-failure is roughly 28 yr.

1 Introduction

The Hubble Space Telescope (HST) uses gyroscopes (gyros) as one layer within its guidance system. Each gyro is packaged in a Rate Sensor Assembly (RSA). These RSAs are packaged in pairs in boxes called Rate Sensor Units (RSUs); it is the RSU that is replaced during a servicing mission, so gyros are replaced in pairs. Each RSU is supported by an Electronic Control Unit (ECU).

Thus, there are six gyros on the HST. It is essential for guidance that there be three working gyros; the other three serve as backups. A gyro can be rendered "non-working" by a failure in the gyro itself, or in the electronics support equipment: both have happened.

Any gyro can replace any other gyro: all triads meet specifications; however, some triads give better guidance than others. For the purposes of the present memo, all triads are regarded as equivalent.

So far, the probability that a given gyro will continue to work has been modeled by the Exponential Probability Law: this is equivalent with a constant failure rate. The purpose of this essay is to show that flex lead failures within gyros show an unmistakable *wear out* behavior, and are therefore better modeled by a Probability Law with a failure rate that is initially vanishingly small, and then increases explosively after some incubation time. A Weibull Probability Law can do this, and is used in this essay. Based on this, we have presented elsewhere new plots for the probability that there will be at least three working gyros in the future.

2 Gyro history

HST was placed into orbit on 25/Apr/90 with six functional gyros: Launch-G1 (SN 108), Launch-G2 (SN 113), Launch-G3 (SN 110), Launch-G4 (SN 138), Launch-G5 (SN 104), and Launch-G6 (SN 127).

The first six gyros: Launch-G1 served until 19/Nov/92, when a capacitor in the input of the power supply in the ECU supporting Launch-G1 failed, de-powering this gyro. The ECU was replaced as part of Servicing Mission 1 (SM1; Dec 93), and this gyro then ran until a flex lead failed on 13/Nov/99. This gyro accumulated 4.63 yr of run-time before failing of a broken flex lead.

Launch-G2 was still running when it was removed as part of SM3a (Dec 99); it accumulated 4.61 yr of run-time.

Launch-G3 was still running when it was removed as part of SM1 (Dec 93); it accumulated 4.88 yr of run-time. It was refurbished and became SM3a-G2, and is still running.

Launch-G4 became unavailable upon the failure of a hybrid circuit in the supporting electronics on 19/Jun/90. It was still running when it was removed as part of SM1 (Dec 93); it accumulated 3.67 yr of run-time. It was refurbished and became SM3a-G3, and is still running.

Launch-G5 was still running when it was removed as part of SM1 (Dec 93); it accumulated 5.03 yr of run-time.

Launch-G6 became unavailable upon the failure of a hybrid circuit (similar to the one that crippled Launch-G4) in the supporting electronics on 19/Jun/90. It continued to run until a flex lead broke on 7/Oct/92, for a run-time of 4.05 yr.

Between SM1 (Dec 93) and SM3a (Dec 99): SM1 (Dec 93) removed the two RSUs containing Launch-G3 (SN 110), Launch-G4 (SN 138), Launch-G5 (SN 104), and Launch-G6 (SN 127). It replaced these with two new RSUs, containing SM1-G3 (SN 112), SM1-G4 (SN 158), SM1-G5 (SN 118), and SM1-G6 (SN 151). SM1 also replaced the ECU supporting Launch-G1, restoring this to working condition. Thus, SM1 placed a full compliment (6) of working gyros in HST.

All six gyros were working up to the Second Servicing Mission (SM2, Feb 97), and no gyros were replaced. However, after SM2, one gyro failed in April 1997, a second failed in 1998, a third failed in early 1999, and a fourth failed in November 1999: this induced a "safe hold" in HST. The planned SM3 was quickly remapped into two parts, SM3a (Dec 99) focusing on replacing gyros to restore HST operation, and SM3b (presently planned for 20/Jan/02) to upgrade instruments. The details for the gyros are as follows:

SM1-G3 failed when a flex lead broke on 20/Ap/99; it accumulated 6.10 yr run-time.

SM1-G4 failed when a flex lead broke on 9/Ap/97; it accumulated 3.60 yr run-time.

SM1-G5 was still running when it was removed as part of SM3a (Dec 99); it accumulated 6.71 yr of run-time.

SM1-G6 failed when a flex lead broke on 28/Oct/98; it accumulated 5.47 yr run-time.

SM3a (Dec 99) until present (Aug 01): SM3a (Dec 99) removed all three RSUs, containing Launch-G1 (SN 110), Launch-G2 (SN 138), SM1-G3 (SN 112), SM1-G4 (SN 158), SM1-G5 (SN 118), and SM1-G6 (SN 151).

It replaced these with three new RSUs, containing SM3a-G1 (SN 155), SM3a-G2 (SN 110) (this is the refurbished Launch-G3), SM3a-G3 (SN 104), SM3a-G4 (SN 138) (this is the refurbished Launch-G4), SM3a-G5 (SN 156), and SM3a-G6 (SN 159). Thus, SM3a placed a full compliment of working gyros in HST.

SM3a-G1 is presently "stored", and is expected to work whenever it is turned on. As of 1/Aug/01, it had run for 0.77 yr.

SM3a-G2, SM3a-G3, and SM3a-G4 are presently working. As of 1/Aug/01, these have run times of 0.51 yr, 2.01 yr, and 1.98 yr, respectively.

SM3a-G5 spun down on 28/Ap/01, probably because a bit of the lubricating material used to support the spinning element during initial operation (once up to speed, this element moves on gas bearings only) assembled itself into a glob so large that it reached the spinning element and imposed excessive friction on it: this is called a "lube patch" failure. This gyro accumulated 1.59 yr run-time.

SM3a-G6 has shown anomalous behavior since shortly after being started: there is an "out of family" large bias drift. This anomalous behavior is not satisfactorily explained. SM3a-G6 is still able to serve as a working gyro; however, its anomalous behavior suggests that its effective lifetime as a working gyro may be limited. Presently, it is "stored", but is expected to work whenever it is turned on. As of 1/Aug/01, it had run for 1.41 yr.

3 The gyro failures

These gyros have three known failure modes. This section review these modes, and describes the probability law governing the likelihood of each mode.

3.1 Failure of supporting electronics

A review of the failure behavior described in the previous section shows that we can dismiss the two failures caused by defective hybrid circuits when we are modeling future behaviors since these hybrids have been replaced by different and much more robust circuits.

We cannot dismiss the possibility of another failure in the ECU. We will model the probability of an ECU failure using an exponential probability law with a mean time to failure (MTTF) of τ_{ECU}.

$$P_{ECU}(t) = \exp(-t/\tau_{ECU}) . \tag{1}$$

This model assigns the constant failure rate $R_{ECU} = 1/\tau_{ECU}$ to this class of failures. An exponential probability law is appropriate for the behavior of shorting behaviors for the kind of capacitor that had shorted within the ECU, and is also appropriate for a variety of other electronics failure modes other than shorting capacitors. We could increase the expected MTTF of ECU failure if we knew that the capacitor that was found to have shorted, was replaced in all cases with a type that would be far less likely to short in this application.

3.2 Lube patch failure

We cannot dismiss the possibility of another "lube patch" failure. Experience with similar gyros used in other programs teaches us that this is a recurring problem with this method of supporting the spinning element. We will model this using an exponential probability law with a mean time to failure (MTTF) of τ_{lube}.

$$P_{lube}(t) = \exp(-t/\tau_{lube})\,. \tag{2}$$

With only one such failure observed among the HST gyros, there is no useful statistical method available to decide whether this probability law is actually appropriate. Nor is the underlying modeling (involving migration of lubrication, among other things) developed far enough to judge whether this probability law is appropriate. We will use it until we have evidence that supports choosing amongst different probability laws.

3.3 Flex lead failure

The dominant mode of failure is by breaking a flex lead. This is known to be caused by corrosion of the coin-silver flex lead by trace materials within the gyro's fill fluid. This corrosion eventually weakens the flex lead until it can no longer support the loads impressed upon it.

While strenuous efforts have been made to decrease the corrosion rate (including attention to sharply reducing the amount of dissolved oxygen and water in the gyro fill fluid, and plating additional silver onto the coin-silver flex leads), it is likely that the gyros now in the HST, as well as the candidate replacements, are still subject to flex lead corrosion. Until events prove otherwise,[1] we must regard this corrosion as the most serious life-limiting process acting on the HST gyros.

The coin silver is an alloy of 85% silver and 15% copper (by weight). The solubility of each of these two metals in the other is limited to less than a percent, in equilibrium, at temperatures below some 200°C. Hence, the

[1] There are established statistical methods for evaluating when changes in construction actually produce better performance. It seems of interest to apply these to the HST gyros. That may be the subject of a future memo.

flex lead material is composed of grains of almost pure silver, surrounded by a matrix of almost pure copper. The alloy resembles a structure built of silver-rich blocks, joined by copper-rich mortar; there are silver-rich filaments joining the blocks. Under the action of the corrosive material, containing bromine and chlorine, the copper-rich mortar is removed from the interstices of the blocks, and converted to a copper-halide deposit, which builds up as a crust on the flex lead; however, some forms within the interstices. A mole of copper halide occupies several times the volume of a mole of copper, so this interstitial buildup exerts an internal pressure that tends to force the silver-rich grains apart: this is equivalent to an applied tensile force. One consequence of this internal pressure is the frequent spalling of the copper halide crust from the outside of the flex lead; another is its eventual rupture.

Also, the copper halide is formed as a fragile assemblage of small grains: this material has negligible tensile strength. Thus, the tensile strength of the flex lead steadily decreases as the copper-rich mortar is corroded into copper halide.

The flex lead, and the design of the gyro, have been chosen so that the initial strength of the flex lead is far greater than the applied loads. Thus, the initial failure rate, caused by a breaking flex lead, is essentially zero, and it must remain zero until corrosion reduces the strength, and increases the internal pressure, until the decreasing strength equals the increasing load.

Inspection of the broken flex leads shows that most of it is not corroded; the corrosion is limited to a band a few millimeters long, starting close to the end of the flex lead where it is soldered to a pin that supplies (or sinks) electrical current. The entire batch of flex lead is carefully prepared to have high strength: this is done by work hardening it. Work hardening distorts the grains, and introduces many tangled dislocations: this inhibits further motion of dislocations, which increases the strength of the flex lead. It also affects the corrosion rate dramatically. The soldering operation, which involves heating the pin and the end of the flex lead to at least 220°C and then letting them cool, is a heat treatment that advances recrystallization, leaving a "heat affected zone" on which the corroding fluid can quickly act. This inadvertent heat treatment is not tightly controlled, and can be expected to be a source of variation among different flex leads, and different gyros.

Cartoon model of flex lead fracture: The following is a sketch of some of the physical principles controlling the breaking of a flex lead. Massive simplifications are made, so the model is only a "cartoon". But even cartoons can be a useful illustration of a complicated situation.

The flex lead will not break until its strength decreases below the applied force, which must be a tensile force (not a compressive one) in order to pull the ends apart. The strength of the uncorroded flex lead is $S = Y \cdot A$, where Y is the yield modulus of the flex lead material, and A is the cross section of the flex lead. It seems apt to assign zero tensile strength to the corrosion product, since this is formed as separate and unconsolidated minute grains. Therefore, the strength of the partially corroded flex lead is, at least approximately,

$$S(t) = Y \cdot < A(t) > \tag{3}$$

where $< A(t) >$ is the effective cross-sectional area of the flex lead; corrosion will reduce this with time.

Corrosion is (at least) a two-step process. First, the corrosive agent and the corrodible material must come together, usually as one of them diffuses to the location of the other. Second, the chemical reaction (corrosion) must happen: this proceeds as a chemical rate process. We would have to attend to each process, diffusion and reaction, if they happened at comparable speeds; however, we will suppose that the diffusion is much slower, and therefore acts as the all-over rate-limiting process. This seems consistent with the slow rate of ionic motion in these metals, and the fast rate of chemical reaction between bromine, chlorine, copper, and silver. Then the fraction of the corroded material within the flex lead is given by the solution of a diffusion equation. Since the flex lead is some 17 times wider than it is thick, the diffusion is dominantly one-dimensional, through the thickness h of the flex lead. We will take this direction as the x-axis, running from $x = 0$ to $x = h$. The concentration $c = c(x,t)$ of the corroded material is described by

$$\frac{\partial c}{\partial t} = -D \cdot \frac{\partial^2 c}{\partial x^2}, \tag{4}$$

where the initial concentration vanishes everywhere: $c(x, t = 0) = 0$ for all x between 0 and h. And where c always remains zero outside the flex lead: $c(x = 0, t) = c(x = h, t) = 0$. Finally, D is the diffusion parameter for the motion of the corrosion into the flex lead.

The solution for the relative concentration of corroded material is

$$c_r(x,t) = 1 - \frac{4}{\pi} \cdot \sum_{j=0}^{\infty} \sin[((2j+1)\pi/h)x] \exp[-((2j+1)\pi/h)^2 Dt] \,. \tag{5}$$

We are interested in the corroded fraction within the flex lead, which is

$$C(t) = \int_0^h c(x,t) = 1 - \frac{8}{\pi^2} \cdot \sum_{j=0}^{\infty} \frac{1}{(2j+1)^2} \exp[-((2j+1)\pi/h)^2 Dt] \,. \tag{6}$$

The non-corroded fraction is $\bar{C} = 1 - C(t)$, and so an estimate for the effective cross-sectional area of the flex lead is $< A(t) >= A \cdot \bar{C}(t)$, and an estimate for the strength of the flex lead is

$$S(t) = Y \cdot A \cdot \bar{C} = Y \cdot A \cdot \frac{8}{\pi^2} \cdot \sum_{j=0}^{\infty} \frac{\exp[-((2j+1)\pi/h)^2 Dt]}{(2j+1)^2} \,; \tag{7}$$

this starts at the value $S(0) = YA$, and decreases smoothly with increasing time, asymptoting to vanishing strength as $t \to \infty$. This expression could be tested by measuring the strength of specimens of flex leads that have been soaked in gyro fluid for different lengths of time.

Equation 7 is complicated at first sight, since it is an infinite series. Inspection shows that the ratio of the $(j = 0)$-term to the $(j = 1)$-term is $9\exp(8\pi^2 Dt/h^2)$, and this rapidly becomes large as Dt/h^2 exceeds unity. Then the entire series is well-approximated by its $(j = 0)$-term. On the other hand, an increasing number of terms is needed as Dt/h^2 approaches zero, and it is then better to recast the series into an alternate form; however, we will not need to consider this case, since we know that the flex lead material and dimensions have been selected so that the initial strength $S(0) = YA$ is far larger than the applied loads. Therefore, we approach failure only when this initial strength is decreased to a small fraction of its initial value, and for this case,

$$S(t) \approx \frac{8YA}{\pi^2} \cdot \exp(-\pi^2 Dt/h^2) \,. \tag{8}$$

The flex lead will fail by breaking at the time t_f when its strength falls to the applied load, say, F_a.

$$t_f \approx \frac{h^2}{\pi^2 D} \cdot \ln\left[\frac{8YA}{\pi^2 F_a}\right] \,. \tag{9}$$

We can give this a "sanity check" as follows. Rearrange Equation 9 to give $D \approx (h^2/\pi^2 < t_f >) \cdot \ln(8YA/\pi^2 F_a)$, use the average time to failure, $< t_f >= 4.77$ yr and the flex lead thickness, $h = 0.60$ mil ($= 15\ \mu$m). This estimates D as $(1.5 \times 10^{-15}$ s$)\cdot \ln(8YA/\pi^2 F_a)$. We now suppose the ratio of the initial strength YA to the applied load F_a is 10, and estimate D as 3.1×10^{-15} s. If we suppose the ratio is 100, and not 10, then we estimate D as 6.6×10^{-15} s. These times have rough, order of magnitude, agreement with diffusion parameters for the motion of copper ions in solid copper.

We might have expected the time to depend on the amount of corrosive material in the gyro fluid. It does not in this cartoon model since we have supposed the chemical reaction times to be short compared with the times for the copper and halogen to come together, and (in effect) that there is always sufficient halogen to reach all the arriving copper to copper halide. If experience shows a halogen-concentration dependence to the corrosion time, then this part of the model must be re-worked.

One source of an applied load is the motion of the float tube during HST maneuvers. However, the design of the leads, and of the entire gyro, has made these loads quite small. Further, as far as we know, these fractures have not happened during an HST maneuver, when we would expect the loads applied to the flex leads to be largest. Therefore, we also consider an alternate source of an applied load, and we find it in the internal pressure caused by the increasing amount of corrosion product accumulating within the flex lead, and around it as a crust. Since the corrosion product is larger than the coin silver precursor, this crust tends to push the flex lead apart, placing a effective tensile load on it. This load will increase (at least roughly) as $C(t)$ does: $F_a(t) = \mathcal{F} \cdot C(t)$ where $\mathcal{F}$ is a parameter adjusted to scale the concentration of corrosion product into the effective applied tensile force; it is the maximum value that this effective force can have, since the maximum value of C is unity.

If we suppose that the corrosion-force parameter is substantially smaller than the initial strength of the flex lead, $\mathcal{F} \ll YA$, then this effect becomes important after a long enough time that we can approximate $C(t)$ using only a single term in it's infinite series:

$$F_a = \mathcal{F} \cdot C(t) \approx \mathcal{F} \cdot [1 - \frac{8}{\pi^2} \exp(-\pi^2 Dt/h^2)]\,. \qquad (10)$$

The flex lead fails when its decreasing strength equals the increasing effective load, at the time:

$$t_f \approx \frac{h^2}{\pi^2 D} \cdot \ln\left[\frac{8YA}{\pi^2 \mathcal{F}} \cdot (1 + \frac{\mathcal{F}}{YA})\right] \approx \frac{h^2}{\pi^2 D} \cdot \ln\left[\frac{8YA}{\pi^2 \mathcal{F}}\right] . \tag{11}$$

This expression is the same as Equation 9, except that the parameter $\mathcal{F}$ has replaced the applied force F_a. Of course, the failure time t_f can also be computed when $\mathcal{F}$ is not much smaller than $\ll YA$, if we find that that is more reasonable.

For both the cases described by Equations 9 and 11, the failure time t_f depends weakly on the ratio of the initial strength of the flex lead to the effective applied force. The failure time t_f depends more strongly on the thickhess h of the flex lead and on the diffusion parameter D.

The diffusion parameter D increases with temperature in an Arrhenius manner: $D(T) = d_0 \exp(-E^\star/RT)$, where d_0 is the high-temperature limit of the diffusion parameter, $E^\star$ is the activation energy for the diffusion process, R is the gas constant, and T is the absolute temperature. The current passing through the flex leads warms the gyro fluid, and this will increase the corrosion rate, and decrease the time to failure t_f. Until we can get an accurate value for $E^\star$, we will suppose $E^\star/R \approx 8\,000$K, which is a typical value for solid state diffusion; this corresponds to a doubling of D, and a halving of t_f, in 8°C. In general, the gyros do not age as rapidly when they are unpowered (and cool) as when they are powered (and hot).

Also, the corrosion product, copper halide, is essentially non-conductive compared with the copper-rich alloy it replaces, and so the electrical resistance of the mortar increases, and with it, the joule heating of the flex lead. Thus, the temperature of the part of the flex lead that is corroding will experience a "runaway" growth in corrosion rate as the current-carrying parts of the flex lead fall below a threshold value.

If the temperature of the flex lead is changing over its use, then the "calendar time" t alone is not a good measure of the extent of the corrosion-degrading strength remaining in the flex lead. If we are given the temperature history of the flex lead, $T(t)$, then we can compute a "T-scaled" time $t^\star$ that better measures the extent of corrosion:

$$t^\star(t) = \int_0^t \exp\left[\frac{E^\star}{R} \cdot \left(\frac{1}{T_0} - \frac{1}{T(t')}\right)\right] dt' \,, \tag{12}$$

where the reference temperature T_0 is chosen for convenience to be the temperature at which the T-scaled time runs at the same rate as the calendar time. If the gyro is switched on for the time t_{on} and we chose the operating temperature as $T_0 = 343\text{K}$ ($= 70°\text{C}$), and the gyro is switched off for the time t_{off} and is then 20°C cooler than T_0, then

$$t^\star(t) = t_{\text{on}} + \exp\left[8\,000\text{K} \cdot \left(\frac{1}{343\text{K}} - \frac{1}{323\text{K}}\right)\right] \cdot t_{\text{off}} = t_{\text{on}} + 0.236 \cdot t_{\text{off}} \,. \tag{13}$$

That is, the gyro "ages" (by corrosion) about four times more slowly when unpowered (and cool) than when powered (and hot), if the temperature difference is 20°C. As a first approximation, we ignore the time the gyro spends unpowered, and count only the powered time in compiling t_f.

So far, this cartoon model implies that the flex lead will not break until the time advances to t_f, and then breaking is certain. However, this cartoon model also shows that the parameters that determine t_f are subject to considerable variations amongst different gyros, and therefore the values of t_f will also vary. One important variable, possibly the most important one on the basis of inspections of the extent of corrosion along each of the flex leads that have been studied, is the value of D. Thus, an ensemble of gyros will show a spread of flex lead breaking times: this time becomes a "random variable" described by a probability law. This probability law will have a failure rate that is vanishing low for small times, when corrosion has not yet had time to act and the flex lead's strength is still many times larger than the (effective) tensile forces. The failure rate will grow to large values only after a threshold time has passed; this threshold time will depend most strongly on D; the log-term is only slowly varying, and will not contribute as much variations.

Summary of model for flex lead fracture In summary, a flex lead will not fail by corrosion until its strength has been degraded down to the applied forces. Since the initial strength of the flex lead is many times that of the applied forces, and the corrosion rate is small, then there can be no corrosion-induced failures for a substantial time. Then, for the first time, it becomes possible for the flex lead to break. Since the applied forces are variable, and especially because the corrosion rate is highly variable, there will be a substantial scatter in the observed failure times. That is, the flex lead failure rate begins at zero, and remains very small until a threshhold time is approached. Individual gyros will have individual effective values for D (mainly caused by variations in the metal alloy's microstructure), and T, and this will introduce a statistical spread of threshold times.

The cartoon model captures aspects of this. When variations in the quantiles are ignored, we find a specific formula for the failure time t_f, and for its dependence on the temperature history $T(t)$, showing that the failure time is mainly sensitive to the diffusion parameter $D(T)$.

4 Fitting models to the data

The lifetimes of the five gyros failing by flex lead corrosion are:

4.63 yr (Launch-G1)	6.10 yr (SM1-G3)
4.05 yr (Launch-G6)	3.60 yr (SM1-G4)
	5.47 yr (SM1-G6).

The average of these lifetimes, with standard deviation, is (4.77 ± 0.81) yr.

The run times of the eleven gyros that have not broken a flex lead are:

4.61 yr (Launch-G2)	6.71 yr (SM1-G5)	0.77 yr (SM3a-G1)
4.88 yr (Launch-G3)		0.51 yr (SM3a-G2)
3.67 yr (Launch-G4)		2.01 yr (SM3a-G3)
5.03 yr (Launch-G5)		1.98 yr (SM3a-G4)
		1.59 yr (SM3a-G5)
		1.41 yr (SM3a-G6).

In both cases, I am assuming that these times include all the pre-launch run time accumulated during pre-launch testing.

Conventional Weibull analysis The cumulative failure fraction versus time is computed from the observed gyro failure times (all for flex lead fractures), and these fractions are adjusted to accommodate the non-failure times using the Leonard Johnson's method, as modified by Drew Auth. A least square fit of these adjusted fractions F_i versus time t_i, using "t_i regressed on F_i, is then carried out to obtain estimates for the Weibull parameters.[2] The result is $\eta = 5.89$ yr, and $\beta = 4.82$, and the correlation coefficient of the fitted line to the data is $r = 0.989$: this is a good fit.

The assignment of uncertainties to this result is a complex topic, and will not be reported here. Several methods give the results that, with 95% confidence, 4.90 yr $< \eta <$ 6.96 yr, and $3.00 < \beta < 12.5$. Since the 95% confidence limit for β is substantially larger than unity — the Exponential Law — then we can conclude with even greater confidence than 95% that these gyro data are inconsistent with an Exponential Law.

Method of Maximum Likelihood The probability of the data, both the failure times $\{t_i^{[f]}\}_{i=1}^{5}$ and the non-failure times $\{t_j^{[\bar{f}]}\}_{j=1}^{11}$, is computed on the basis of assumed values for the Weibull parameters $\mathbf{p} = (\eta, \beta)$:

$$\mathcal{L} = \prod_{i=1}^{5} f(t_i^{[f]}|\mathbf{p}) \cdot \prod_{j=1}^{11} P(t_i^{[\bar{f}]}|\mathbf{p}) , \tag{14}$$

this is called the "likelihood of the parameters $\mathbf{p}$, given the data". The values of the parameters are varied until this likelihood is maximized, and the resulting values are called the "maximum likelihood estimates" of the parameters.

Experience shows that these two methods are competitive in producing reliable estimates. The results in our case are $\eta = 6.02$ yr, and $\beta = 5.12$: the difference between these results and the former ones is much smaller than those expected from statistical fluctuations.

[2] This procedure is reported as "standard" in various texts including "The New Weibull Handbook" by R. B Abernethy. It is part of modern computer-based Weibull analysis programs.

Final results HST gyro failures are dominated by flex lead corrosion, and this is well-described using a Weibull Law. However, electronics failures and lube patch failures have also happened, and cannot be excluded on the basis of current knowledge. We do not have sufficient failures of these types to establish a failure law for each, but we can estimate it — until more becomes known — as an Exponential Law with $\tau \approx 28$ yr, based on the number of these failures and on the total number of operating hours.

It is plausible that these failure modes operate independently, and therefore the total probability law is the product of the Weibull and the Exponential Laws:

$$P(t) = P_W(t|5.89\text{ yr}, 4.82) \cdot \exp[-t/28\text{ yr}]\,. \tag{15}$$

A General rules for probability and reliability

We begin by stating the probability concepts for a single device whose failure time t_f is a random variable; that is, t_f is given by a probability law. Then the Weibull law is presented, and extended to apply to a set of devices which are preconditioned by a "burn in".

Then the extentions are made to the situation where one operating device is required, and a failing device is replaced with a similar device with $N = 1$, 2, and 3 replacements available.

Finally, the extensions are made to situations where several operating devices are required, and there are N replacements available.

A.1 A single device

Suppose we have a device that either works from the time it is started ($t = 0$) until a time t, or fails at a time t_f. The probability density of a failure at t is denoted $f(t)$; thus, the probability of a failure between $t = 0$ and t, called the "cumulative failure probability", is

$$F(t) = \int_0^t f(t')\,dt', \tag{16}$$

and the probability of no failure over this interval is

$$P(t) = 1 - F(t) = 1 - \int_0^t f(t')\,dt' = \int_t^\infty f(t')\,dt'; \tag{17}$$

the rate of failures is

$$R(t) = -d\ln(P)/dt. \tag{18}$$

Each of these can be computed from any other. For example, $f(t) = dF/dt = -dP/dt$, and $P(t) = \exp[-\int_0^t R(t')\,dt']$.

The Weibull probability Law is

$$\begin{aligned} f(t|\eta,\beta) &= (\beta/\eta)\cdot(t/\eta)^{\beta-1}\cdot\exp[-(t/\eta)^\beta] && (19)\\ F(t|\eta,\beta) &= 1-\exp[-(t/\eta)^\beta] && (20)\\ P(t|\eta,\beta) &= \exp[-(t/\eta)^\beta] && (21)\\ R(t|\eta,\beta) &= (\beta/\eta)\cdot(t/\eta)^{\beta-1}; && (22) \end{aligned}$$

any of these can be used as a definition since each implies all the others.

The Exponential Probability Law is a special case of the Weibull Law, with $\beta = 1$.

Initial run times before start of service Suppose an ensemble of devices are operated for a time s before the devices are actually placed into service. Suppose all the failed devices are removed, and only the surviving devices are actually forwarded into service, and suppose we now start the clock at $t = 0$, and require the probability of failure (or non-failure).

The failure rate aged from its zero-time rate of $R(0)$ to $R(s)$ during the initial run time of duration s, and it will age to $R(t+s)$ after the devices are placed into service. The probability that a device will not fail between $t = 0$ and t is

$$P(t) = \exp[-((t+s)/\eta)^\beta + (s/\eta)^\beta]; \tag{23}$$

note that this incorporates the two things we know: first, the failure rate begins aging as $R(t + s)$ when these devices are placed into service, and *only* operating devices are placed into service: thus, $P(t = 0) = 1$. This expression for $P(t)$ is similar to, but not identical with, the so-called "three parameter Weibull":

$$P(t) = \exp[-((t+s)/\eta)^{\beta}]. \tag{24}$$

A.2 One device has to work, and there are replacements

The probability that a single device, 1, will work over the interval from $t = 0$ to t is $P_1(t)$.

The probability that that device will fail within the small interval $t', t' + \Delta t$ is $f_1(t')\Delta t$. The probability that the replacement, 2, will work from t' until t is $P_2(t-t')$, if we assume that it has not aged while waiting to serve. These are statistically independent events, and so the probability that we start with device 1, and finish with an operating device 2, (recalling that t' is any time between 0 and t) is

$$P_{1\to 2}(t) = \int_0^t f_1(t')P_2(t-t')\,dt'. \tag{25}$$

Applying similar reasoning, the probability that we start with device 1, and then require device 2, and finally finish with device 3, is

$$P_{1\to 2\to 3}(t) = \int_0^t K_{1\to 2}(t'')P_3(t-t')\,dt', \tag{26}$$

where the kernel K is

$$K_{1\to 2}(t) = \int_0^t f_1(t')f_2(t-t')\,dt'. \tag{27}$$

The pattern is now clear: we include a factor $f_i(t')$ for each device i that fails at an intermediate time t_i, and a final factor $P_j(t - t_j)$ for the device j that finishes serving from t_j until t (t_j is the same time as the final device to fail, at $t_{i'}$); we integrate over all the intermediate failure times, $t_1, t_2, \ldots, t_{i'} = t_j$.

It can be useful to denote the sequence of failures and final operating parts using a diagram:

```
(1)-------->   the device 1 operates from t=0 to t

(1)--!(2)-->   the device 1 fails at t', "!",
               and 2 operates from t' to t.
```

It is easy to write down such a diagram, and it translates easily into the integral; it is also easy to reverse the procedure, so these expressions are equivalent. However, the diagrams are more quickly comprehensible.

The total probability of having a single working part is (ignoring the time to switch in the new part, and the reliability of this process), the sum of the separate probabilities, since these represent logically exclusive events:

$$P(t) = P_1(t) + P_{12}(t) + \cdots + P_{12\ldots N}(t)\,, \tag{28}$$

where N is the number of parts available for this system.

These expressions can be computed explicitly when P is the Exponential Probability Law. Then we quickly find $P(t) = [1+(\lambda t)+(\lambda t)^2/2+(\lambda t)^3/3!+\cdots+(\lambda t)^N/N!]\cdot\exp(-\lambda t)$. Since the sequence in square braces is precisely the expansion of $\exp(+\lambda t)$, then the $N \to \infty$ limit is $P(t) = 1$ for all t: the probability of having a working part is unity when there are indefinitely many replacements available.

These expressions cannot be computed explicitly when P is the Weibull Probability Law, and $\beta \approx 4.82$. We must turn to numerical integration, which is straightforward and even simple when we only have to consider $N = 1$; however, it is increasingly tedious for larger values of N.

A.3 Two devices have to work, and there are replacements

The first case is that both initially operating devices continue to work until t:

```
Case 1:  (1)-------->
         (2)-------->
```

We have $P_{c1}(t) = P_1(t) \cdot P_2(t),$.

The second case is that either 1 fails and is replaced by 3, or 2 fails and is replaced by 3.

```
Case 2:  (1)--!(3)-->    or    (1)-------->
         (2)-------->          (2)--!(3)-->
```

These are distinct, so we add their individual probabilities: $P_{c2}(t) = P_{1\to3}(t) \cdot P_2(t) + P_1(t) \cdot P_{2\to3}(t)$.

The above scheme represents the decision to use device 3 to replace either of device 1 or device 2; an alternative scheme is to always replace a failed 1 with 3, and to always replace a failed 2 with 4: the diagrams and integrations are straightforward to write down for this alternative scheme.

The third case involves two failures; this can happen in two different ways, called case 3a and case 3b:

```
Case 3a:  (1)--!(3)--!(4)-->    or    (1)--------------->
          (2)--------------->         (2)--!(3)--!(4)-->
```

and

```
Case 3b:  (1)--!(3)-->
          (2)--!(4)-->
```

The probabilities for these are $P_{\mathrm{c3a}}(t) = P_{1\to3\to4}(t) \cdot P_2(t) + P_1(t) \cdot P_{2\to3\to4}(t)$ and $P_{\mathrm{c3b}}(t) = P_{1\to3} \cdot P_2(t) + P_1(t) \cdot P_{2\to4}(t)$, and the total probability of case 3 is the sum of the two sub-parts: $P_{\mathrm{c3}}(t) = P_{\mathrm{c3a}}(t) + P_{\mathrm{c3b}}(t)$. Tacit within these rules is the decision to use 3 to replace 1, and to use 4 to replace 2.

We can imagine a different replacement strategy: we can replace the device that fails first with the most robust of the remaining devices. This "couples" the integrations: the changes are straightforward and involve integrations over all the failure times, and explicit separations into the subregions corresponding to first failure, and then other failures. We will not show these results here.

These formulae can all be computed explicitly for the Exponential Probability Law; however, only numerical integrations are possible when $\beta \approx 4.82$.

A.4 M devices have to work, and there are replacements

There is nothing new in principle in $N > 2$: all the essential ideas are present when $N = 1$ and $N = 2$.